arXiv:1511.01400v1 [stat.ME] 4 Nov 2015

Multiple Testing with Heterogeneous Multinomial Distributions

Joshua Habiger[*],
David Watts[†], and
Michael Anderson [‡]

November 5, 2015

Abstract

False discovery rate (FDR) procedures provide misleading inference when testing multiple null hypotheses with heterogeneous multinomial data. For example, in the motivating study the goal is to identify species of bacteria near the roots of wheat plants (rhizobacteria) that are associated with productivity, but standard procedures discover the most abundant species even when the association is weak or negligible, and fail to discover strong associations when species are not abundant. Consequently, a list of abundant species is produced by the multiple testing procedure even though the goal was to provide a list of producitivity-associated species. This paper provides an FDR method based on a mixture of multinomial distributions and shows that it tends to discover more non-negligible effects and fewer negligible effects when the data are heterogeneous across tests. The proposed method and competing methods are applied to the motivating data. The new method identifies more species that are strongly associated with productivity and identifies fewer species that are weakly associated with productivity.

Keywords: False Discovery Rate; Heterogeneity; Multinomial; Multiple hypothesis testing; Rhizosphere

1 Introduction

Bacteria that inhabit the interface between the root and the soil (the rhizosphere) play a crucial role in the growth and development of plants. The living portion of the rhizosphere is very complex: it is estimated to contain billions

[*]Department of Statistics, Oklahoma State University, email:jhabige@okstate.edu
[†]Department of Statistics, Oklahoma State University, email: david.watts@okstate.edu
[‡]Department of Plant and Soil Science, Oklahoma State University, email: michael.anderson@okstate.edu

of individual organisms (Gans et al., 2005) and tens of thousands of species of bacteria (rhizobacteria) per gram of soil (Mendes et al., 2013). Some rhizobacteria positively impact plant health by processing soil nutrients into plant available forms (Vessey, 2003), helping plants tolerate stresses (Yang et al., 2009), protecting plants against disease causing organisms (Bakker et al., 2013) and improving soil structure (Alami et al., 2000). Others have a negative impact on plant growth (Suslow and Schroth, 1982). The identification of these rhizobacteria will allow for 1) the evaluation of current agricultural practices, 2) a more rigorous definition of productive and healthy soil systems and 3) the development of technologies to manipulate the rhizosphere for greater agricultural productivity and sustainability.

The goal in Anderson and Habiger (2012) was to identify/discover species of rhizobacteria in wheat whose abundance is associated with productivity. The statistical problem, formally described in Section 2 of this manuscript, amounts to testing multiple null hypotheses simultaneously with multinomial data. Standard statistical methods, outlined in Section 3, compute a test statistic for each null hypothesis and apply a false discovery rate (FDR) multiple testing procedure to the collection of test statistics. For example, the well-known BH procedure (Benjamini and Hochberg, 1995) or an adaptive BH procedure (Storey et al., 2004; Liang and Nettleton, 2012) can be applied to the p-values or the local FDR procedure, introduced in Efron et al. (2001) and formalized in Sun and Cai (2007), could be applied to the Z-scores for the tests.

In Sections 4 and 7 we see that these standard approaches provide misleading inference. In particular, some rhizobacteria that appear to be strongly associated with productivity are not discovered while other rhizobacteria that are only negligibly associated with productivity are discovered. A more careful inspection reveals that standard methods will, in general, tend to discover the most abundant rhizobacteria, even if they are only negligibly associated with productivity. Consequently, some abundant rhizobacteria are mislabeled as productivity-associated rhizobacteria and many productivity-associated bacteria are not discovered because they are not abundant. This confounding inference can have a significant detrimental impact on efforts to understand how the rhizosphere interacts with the environment. A similar phenomenon was observed in Sun and McLain (2012), where the goal was to identify high schools whose test scores were associated with socio-economic status, but standard approaches tended to identify the largest schools instead.

It turns out that standard methods fail because the sample sizes for the multinomial data are heterogeneous across tests. Section 5 provides a multiple testing procedure for multinomial data which incorporates this heterogeneity via conditional local FDR (clFDR) statistics. The statistics are referred to as "conditional" to emphasize the fact that analysis is based on a multinomial model for each test, which arises by conditioning on the (random) sample sizes. The result is a procedure that depends functionally upon the sample sizes across tests that, as we will show, controls the FDR. Section 6 shows analytically that the clFDR procedure provides more scientifically relevant inference and shows that the improvement can be characterized as a "thresholding effect". That

Table 1: Depiction of the data in Anderson and Habiger (2012). Shoot biomass x_n in grams for groups $n = 1, 2, ..., 5$ was 0.86, 1.34, 1.81, 2.37, and 3.00, respectively. Row totals are in the last column.

Species m	y_{1m}	y_{2m}	y_{3m}	y_{4m}	y_{5m}	Total (n_m)
1	0	1	1	0	5	7
2	9	2	0	0	3	14
⋮	⋮	⋮	⋮	⋮	⋮	⋮
778	16	10	29	18	13	81

is, the rejection region for a test based on the clFDR statistic is decreased when the sample size is large and increased when the sample size is small. Consequently, the probability of discovering a negligible association is decreased and the probability of discovering a strong association is increased. Section 7 applies the new method to the motivating data and verifies that it performs as anticipated.

2 Motivating data and objective

The basic objective is to identify rhizobacteria that are *significantly* associated with wheat shoot biomass, which is a measure of wheat productivity. For the moment we use the term "significantly" somewhat freely. Clarification is provided in Section 4.

The motivating data are depicted in Table 1. For a detailed account of the experiment see Anderson and Habiger (2012). Here, y_{nm} represents the number of DNA copies, or abundance, of the mth rhizobacterial species among the nth group of wheat plants for $n = 1, 2, ..., N$ and $m = 1, 2, ..., M$. For this data $N = 5$ and $M = 778$. The average shoot biomass among wheat plants in the nth group is denoted x_n and the vector of biomass measurements is denoted by $\mathbf{x} = (x_1, x_2, ..., x_N)^T$. Denote the vector of abundance measurements for species m by $\mathbf{y}_m = (y_{1m}, y_{2m}, ..., y_{Nm})^T$ and the random vector by $\mathbf{Y}_m$. Denote the $M \times N$ random matrix by $\mathbf{Y} = (\mathbf{Y}_1^T, \mathbf{Y}_2^T, ..., \mathbf{Y}_M^T)^T$. Denote the random sample size (row total in Table 1) for $\mathbf{Y}_m$ by $N_m = \sum_{n=1}^{N} Y_{nm}$ and a realization of N_m by n_m.

While more complex models could be considered, here we consider a log-linear model for each species for simplicity and to facilitate parameter estimation and mixture modeling later. Specifically, assume Y_{nm} has a Poisson distribution with mean μ_{nm} and assume that $log(\mu_{nm}) = \alpha_m + \beta_m x_n$, where each α_m and β_m are regression parameters taking values in $\Re$. Further assume $Y_{1m}, Y_{2m}, ..., Y_{Nm}$ are independent for each m. Observe that if $\beta_m = 0$ then $\mu_{nm} = \alpha_m$ for each x_n, and hence the abundance of species m is not associated with productivity. If β_m is positive/negative then μ_{nm} is increasing/decreasing in x_n (note that

$x_1 < x_2 < ... < x_N$). Thus, productivity-associated bacteria can be identified by testing the null hypothesis $H_m : \beta_m = 0$ for each of the $M = 778$ bacteria. The unobservable state of a null hypothesis is denoted by $\theta_m = I(\beta_m \neq 0)$, where $I(\cdot)$ is the indicator function. The decision to reject or retain H_m using $\mathbf{Y}$ is denoted $\delta_m(\mathbf{Y}) = I(H_m \text{ rejected})$.

A p-value or Z-score for H_m can be based on sufficient statistic $T_m = \mathbf{x}^T\mathbf{Y}_m$. To avoid estimation of the nuisance parameter α_m, tests based on this log-linear model often utilize the conditional distribution of $\mathbf{Y}_m|N_m = n_m$, which has multinomial probability mass function

$$p(\mathbf{y}_m|n_m;\beta_m) = \frac{n_m!}{\prod_{n=1}^N y_{nm}!}\prod_{n=1}^N p_n(\beta_m)^{y_{nm}} \text{ where } p_n(\beta_m) = \frac{\exp\{\beta_m x_n\}}{\sum_{n=1}^N \exp\{\beta_m x_n\}}. \tag{1}$$

For details and additional motivation see McCullagh and Nelder (1989). Denote the multinomial probability vector by $\mathbf{p}(\beta_m) = [p_1(\beta_m), p_2(\beta_m), ..., p_N(\beta_m)]^T$. Now, observe that
$E[T_m|n_m;\beta_m] = n_m\mathbf{x}^T\mathbf{p}(\beta_m)$ where $E[T_m|n_m;\beta_m]$ denotes the expectation of T_m taken with respect to the probability mass function in (1), and $Var(T_m|n_m;\beta_m) = n_m\mathbf{x}^T\mathbf{\Sigma}(\beta_m)\mathbf{x}$ where $\mathbf{\Sigma}(\beta_m) = [diag(\mathbf{p}(\beta_m)) - \mathbf{p}(\beta_m)\mathbf{p}(\beta_m)^T]$. Thus, under $H_m : \beta_m = 0$ a Z-score can be computed

$$Z_m = \frac{T_m - n_m\mathbf{x}^T\mathbf{p}(0)}{\sqrt{n_m\mathbf{x}^T\mathbf{\Sigma}(0)\mathbf{x}}}.$$

A p-value for Z_m can be computed $P_m = 2[1 - F_0(|Z_m|)]$, where $F_0(z)$ is the cumulative distribution function of Z_m under Model (1) when $\beta_m = 0$. Note that F_0 can be approximated with a standard normal distribution function or the exact distribution of F_0 can be simulated by sampling from $p(\mathbf{y}_m|n_m;\beta_m)$ with $\beta_m = 0$.

Typically a multiple decision function $\delta(\mathbf{Y}) = [\delta_1(\mathbf{Y}), \delta_2(\mathbf{Y}), ..., \delta_M(\mathbf{Y})]^T$ which controls the FDR at some fixed level α is employed in this "large M" setting. To define the FDR, denote the number of discoveries by $R(\mathbf{Y}) = \sum_{m=1}^M \delta_m(\mathbf{Y})$ and denote the number of false discoveries by $V(\theta,\mathbf{Y}) = \sum_{m=1}^M \delta_m(\mathbf{Y})(1-\theta_m)$, where $\theta = (\theta_1, \theta_2, ..., \theta_M)^T$. Then

$$FDR = E\left[\frac{V(\theta,\mathbf{Y})}{R(\mathbf{Y})}\Big| R(\mathbf{Y}) > 0\right]\Pr(R(\mathbf{Y}) > 0)$$

where the expectation is taken over $\mathbf{Y}$. See Benjamini and Hochberg (1995) or Storey (2003) for alternative definitions and discussions.

3 Standard FDR procedures

Standard FDR procedures are defined in terms of the Z-scores or p-values. For example, the well-known Benjamini and Hochberg (1995) procedure ranks the p-values $P_{(1)} \leq P_{(2)} \leq ... \leq P_{(M)}$ and rejects the k null hypotheses with the

smallest p-values, where $k = \max\{m : P_{(m)} \leq \alpha m/M\}$. If $P_{(m)} > \alpha m/M$ for each m then no null hypotheses are rejected. Adaptive p-value procedures (Storey et al., 2004; Liang and Nettleton, 2012) for FDR control operate in a similar fashion, but incorporate an estimate of the proportion of true null hypotheses. See Blanchard and Roquain (2009) for a more comprehensive list and comparison of p-value procedures for FDR control.

Local FDR procedures based on Z-scores typically utilize a random mixture model, perhaps first considered within the context of the FDR in Efron et al. (2001); Genovese and Wasserman (2002); Storey (2003). For additional references see Efron (2010). For example, assume $Z_1, Z_2, ..., Z_M$ are independent and identically distributed with mixture density $f(z)$ defined by

$$f(z) = \pi_0 f_0(z) + (1 - \pi_0) f_1(z),$$

where $\pi_0 \in (0, 1)$ is the mixing proportion and f_0 and f_1 are the densities when H_m is true and false, respectively. If the state of H_m is viewed as random then π_0 is the prior probability that H_m is true, i.e. $\pi_0 = \Pr(\theta_m = 0)$. The local FDR under mixture density f is defined by

$$lFDR(z; f) = \frac{\pi_0 f_0(z)}{f(z)}$$

and the local FDR statistic for Z_m is defined $lFDR_m = lFDR(Z_m; f)$.

The lFDR procedure in Sun and Cai (2007) is operationally implemented as follows. First, rank the lFDR statistics $lFDR_{(1)} \leq lFDR_{(2)} \leq ... \leq lFDR_{(M)}$. If $\sum_{i=1}^{m} lFDR_{(i)} > m\alpha$ for each m then no null hypotheses are rejected. Otherwise, the k null hypotheses corresponding to $lFDR_{(1)}, lFDR_{(2)}, ..., lFDR_{(k)}$ are rejected, where

$$k = \max\left\{ m : \sum_{i=1}^{m} lFDR_{(i)} \leq m\alpha \right\}.$$

The lFDR procedure can also be written $\delta_m(\mathbf{Z}) = I(lFDR_m \leq \lambda)$, where

$$\lambda = \begin{cases} 0 & \text{if k = 0} \\ lFDR_{(k)} & \text{otherwise.} \end{cases} \tag{2}$$

While this thresholding notation seems redundant, it will be useful for studying lFDR procedures later. Sun and Cai (2007) showed that local FDR procedures dominate FDR methods based on p-value statistics in that among all procedures with asymptotic $FDR \leq \alpha$, they have the smallest missed discovery rate, defined

$$MDR = \frac{E[S(\theta, \mathbf{Y})]}{E[M - R(\mathbf{Y})]}$$

where $S(\theta, \mathbf{Y}) = \sum_{m=1}^{M} \theta_m[1 - \delta_m(\mathbf{Y})]$ is the number of erroneously retained null hypotheses and $M - R(\mathbf{Y})$ is the number of retained null hypotheses. Hence we focus on lFDR procedures for the remainder of this manuscript.

Table 2: Parameter estimates, the AIC and BIC, and the number of discoveries $R(\mathbf{Y})$ when implementing the adaptive $lFDR$ procedure at $\alpha = 0.05$ are summarized when using a 3, 4, and 5 component normal mixture density, with one component density being the standard normal density.

$\hat{\pi}_0, (\hat{\pi}_1, \hat{\mu}_1, \hat{\sigma}_1^2), (\hat{\pi}_2, \hat{\mu}_2, \hat{\sigma}_2^2), \ldots$	AIC	BIC	$R(\mathbf{Y})$
0.68, (0.22,-1.5,3.3^2), (0.10,1.9,0.9^2)	3139.6	3167.6	85
0.21, (0.01,-9.0,5.7^2), (0.11, -1.3, 2.6^2), (0.67,0.1,1.6^2)	3082.2	3124.1	175
0.36, (0.2,-6.8,5.5^2), (0.53,-0.3,1.8^2), (0.07,1.1,0.5^2), (0.01, 2.7,0.3^2)	3085.0	3141.0	125

In practice f is not known but can be consistently estimated, thereby rendering the resulting lFDR procedure adaptive. To illustrate, assume f is a mixture of 3, 4, or 5 normal densities with one of the component densities being a standard normal (null) density. Maximum likelihood estimates (found using the mixtools package in R), the AIC and BIC for each model, and the number of discoveries made by the adaptive lFDR procedure are summarized in Table 2.

4 Non-negligible effects

So which model should we use? Observe that the 4 and 5 component mixture models could be deemed preferable over the 3 component model because they have lower AIC and BIC values and lead to more discoveries. However, each of these models have a non-null component density that is very near the null standard normal density. For example, one of the densities in the 4 component model has mean 0.1 and variance 1.6^2 while one component density has mean -0.3 and variance 1.8^2 for the 5 component model. If the objective of the analysis is to discover *non-negligible* effects, we may be tempted to reconsider these two models because a discovery could mean that the Z-score was generated from a density that is only negligibly different from the null density.

However, a more careful inspection reveals that a species with a Z-score from a "near-null" component density might be more scientifically meaningful than Z-scores from a component density that is farther from the null and vice versa. To see this, for $\beta_m = \gamma > 0$ and $N_m = n$, denote the conditional mean of Z_m under model (1) by

$$\mu(n, \gamma) \equiv E[Z_m | N_m = n; \beta_m = \gamma] = \frac{\sqrt{n}\mathbf{x}^T[\mathbf{p}(\gamma) - \mathbf{p}(0)]}{\sqrt{\mathbf{x}^T \boldsymbol{\Sigma}(0)\mathbf{x}}} \tag{3}$$

and the conditional variance by

$$\sigma^2(\gamma) \equiv Var(Z_m | N_m = n; \beta_m = \gamma) = \frac{\mathbf{x}^T \boldsymbol{\Sigma}(\gamma)\mathbf{x}}{\mathbf{x}^T \boldsymbol{\Sigma}(0)\mathbf{x}}. \tag{4}$$

Note that $\mathbf{p}(0.1)^T = (0.18, 0.19, 0.20, 0.21, 0.22)$ deviates from the null probabilities $\mathbf{p}(0) = \mathbf{0.2}$ by a small, perhaps negligible, amount. When $\gamma = 0.3$ or $\gamma = 0.5$, which we refer to as moderate effects and large effects, respectively, probabilities are farther from $\mathbf{0.2}$. For example, $\mathbf{p}(0.5)^T = (0.11, 0.14, 0.18, 0.24, 0.33)$. Figure 1 plots $\mu(n, \gamma)$ vs. n for $\gamma = 0.1, 0.3, 0.5$ and displays the distribution of the n_ms for the data in Table 1. Observe that even for moderate effects, Z_m has mean between 0 and 1 when n_m is less than or equal to 10 and $n_m \leq 10$ for the majority (59%) of tests. On the other hand, when n is large $\mu(n, \gamma)$ can be large even for negligible effects, and n can be as large as 911. For example, $\mu(911, 0.1) = 2.28$. For moderate effects, $\mu(911, 0.3) = 6.86$. Hence, we may anticipate a procedure based on Z-scores alone to reject H_m when n_m is large even when β_m is only negligibly different from 0. Likewise, H_m will likely be retained even when β_m is significantly different (in a practical sense) from 0 when n_m is small. For this particular data set, this means that species which are merely the most abundant (large n_m) will tend to be identified as productivity-associated even if the association is weak or negligible, while species that are strongly correlated with productivity may not be discovered because they are not abundant.

The fundamental issue is that any procedure based on Z-scores (or p-values) alone can only be used to detect large $|\mu(n_m, \beta_m)|$, which may or may not indicate that the parameter of interest β_m is negligibly different from 0. Indeed, for *any* $\beta_m \neq 0$, $|\mu(n_m, \beta_m|$ is monotonically increasing in n_m. Hence for large enough n_m we should anticipate $|Z_m|$ to be large even if β_m is only negligibly different from 0. Next we develop an FDR procedure which allows for more direct inference on β_m.

5 A conditional local FDR procedure

The proposed procedure is based on a mixture of multinomial distributions. We first define the procedure assuming parameters in the model are known. Parameter estimation follows.

5.1 The procedure

Assume that $\mathbf{Y}_m | N_m = n_m$ is generated from a mixture of multinomial probability mass functions. Specifically, let $\gamma = (\gamma_0, \gamma_1, ..., \gamma_K)^T$ be the collection of $K+1$ distinct possible values for β_m in $\Re$ with $K \geq 1$ and $\gamma_0 = 0$. Denote mixing proportions by $\pi = (\pi_0, \pi_1, ..., \pi_K)^T$ satisfying $\sum_{k=0}^{K} \pi_k = 1$ and $0 < \pi_k < 1$ for each k. Define the mixture probability mass function

$$p(\mathbf{y}_m | n_m; \gamma, \pi) = \sum_{k=0}^{K} \pi_k p(\mathbf{y}_m | n_m; \gamma_k) = \sum_{k=0}^{K} \pi_k \frac{n_m!}{\prod_{n=1}^{N} y_{nm}!} \prod_{n=1}^{N} p_n(\gamma_k)^{y_{nm}} \quad (5)$$

where $p_n(\gamma_k)$ is the multinomial probability from the log-linear model defined as in (1). Denote $\theta_m = I(\beta_m \neq \gamma_0)$ so that the mth null hypothesis is $H_m : \theta_m = 0$

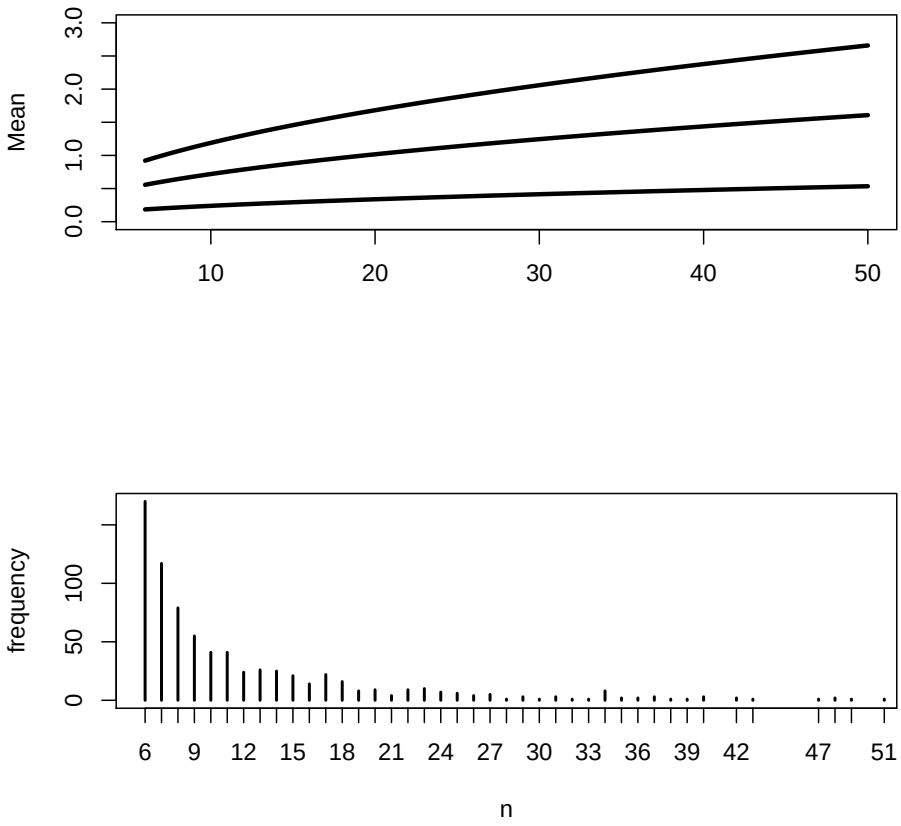

Figure 1: Top: Plot of $\mu(n, \gamma)$ vs. n for $\gamma = 0.1, 0.3, 0.5$ from bottom to top, respectively. Bottom: Bar chart of n_ms. The n_ms greater than 51 (27 of them) are not depicted.

or equivalently $H_m : \beta_m = \gamma_0$. Observe that π_0 is still the proportion of true null hypotheses or the prior probability that H_m is true.

Given parameters π, γ, define

$$clFDR(\mathbf{y}_m, n_m; \pi, \gamma) = \frac{\pi_0 p(\mathbf{y}_m | n_m; \gamma_0)}{p(\mathbf{y}_m | n_m; \gamma, \pi)}$$

and define the conditional lFDR statistic for H_m by $clFDR_m = clFDR(\mathbf{Y}_m, N_m; \pi, \gamma)$. To define the clFDR procedure, first rank the clFDR statistics $clFDR_{(1)} \leq clFDR_{(2)} \leq ... \leq clFDR_{(M)}$. If $\sum_{i=1}^{m} clFDR_{(i)} > m\alpha$ for each m, reject $k = 0$ null hypotheses. Otherwise, reject the $k = \max\{m : \sum_{=1}^{m} clFDR_{(i)} \leq m\alpha\}$ null hypotheses corresponding to $clFDR_{(1)}, clFDR_{(2)}, ..., clFDR_{(k)}$. Formally, the mth clFDR decision function is defined $\delta_m^c(\mathbf{Y}) = I(clFDR_m \leq \lambda)$, where

$$\lambda = \begin{cases} 0 & \text{if k = 0} \\ clFDR_{(k)} & \text{otherwise.} \end{cases}$$

and denote the clFDR multiple decision function by $\delta^c(\mathbf{Y}) = [\delta_1^c(\mathbf{Y}), \delta_2^c(\mathbf{Y}), ..., \delta_M^c(\mathbf{Y})]^T$.

Under the model in (5) the clFDR procedure controls the FDR. A proof is formally presented below for completeness, but is essentially the same as the proof of Theorem 1 in Cai and Sun (2009) with a few minor notational modifications. The important point is that the proof carries over as long as $clFDR_m$ is the posterior probability that the null hypothesis is true, i.e. $clFDR_m = \Pr(\theta_m = 0 | \mathbf{y}_m, n_m; \pi, \gamma)$.

Theorem 1 *For $0 < \alpha \leq 1$, $\delta^c(\boldsymbol{Y})$ has $FDR \leq \alpha$ under the model in (5).*

Proof: Observe that given $\boldsymbol{Y}_m = \boldsymbol{y}_m$ and $N_m = n_m$, $clFDR_m = \Pr(\theta_m = 0 | \boldsymbol{y}_m, n_m; \pi, \gamma)$ and that if $R(\boldsymbol{Y}) = k > 0$ then $E[V(\theta, \boldsymbol{Y}) | \boldsymbol{Y}, R(\boldsymbol{Y}) = k] = \sum_{i=1}^{k} clFDR_{(i)} \leq k\alpha$ by construction. Thus, by the law of iterated expectation we have

$$\begin{aligned} FDR &= E\left[\sum_{k=1}^{M} \frac{1}{k} E\left[V(\theta, \boldsymbol{Y}) | \boldsymbol{Y}, R(\boldsymbol{Y}) = k\right] P(R(\boldsymbol{Y}) = k)\right] \\ &\leq E\left[\sum_{k=1}^{M} \frac{\alpha k}{k} \Pr(R(\boldsymbol{Y}) = k)\right] \leq \alpha. \end{aligned}$$

As in the previous section, parameters must be estimated to implement the clFDR procedure. That is, the proposed adaptive clFDR procedure is the same as the clFDR procedure, except that it uses adaptive clFDR statistics, defined $\widehat{clFDR}_m = clFDR(\mathbf{y}_m, n_m; \hat{\pi}, \hat{\gamma})$ for each m where $\hat{\pi}$ and $\hat{\gamma}$ are estimates of π and γ. Maximum likelihood estimation is discussed next.

5.2 Maximum likelihood estimation

Denote $\mathbf{n} = (n_1, n_2, ..., n_M)^T$ and assume that given $\mathbf{n}$, $\mathbf{Y}$ has joint probability mass function

$$p(\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_M | \mathbf{n}; \gamma, \pi) \equiv \prod_{m=1}^{M} p(\mathbf{y}_m | n_m; \gamma, \pi) = \prod_{m=1}^{M} \sum_{k=0}^{K} \pi_k p(\mathbf{y}_m | n_m; \gamma_k).$$

Define log likelihood function

$$l(\gamma, \pi) \equiv \sum_{m=1}^{M} \log \left(\sum_{k=0}^{K} \pi_k p(\mathbf{y}_m | n_m; \gamma_k) \right).$$

To facilitate the EM algorithm (Dempster et al., 1977), let $\mathbf{z}$ be a $(K+1) \times M$ matrix with z_{km} taking on value 1 if $\mathbf{y}_m$ has component pmf $p(\mathbf{y}_m | n_m; \gamma_k)$ and 0 otherwise. Then the complete data log likelihood is

$$l_c(\gamma, \pi) = \sum_{m=1}^{M} \sum_{k=0}^{K} z_{km} \log(\pi_k p(\mathbf{y}_m | n_m; \gamma_k)).$$

Now, let γ^{old} and π^{old} denote current parameter estimates of γ and π, respectively. Then the E and M steps below yield new parameter estimates, denoted γ^{new} and π^{new}.

E: For $k = 0, 1, ..., K$ and $m = 1, 2, ..., M$ compute

$$\hat{z}_{km} = \frac{\pi_k^{old} p(\mathbf{y}_m | n_m; \gamma_k^{old})}{\sum_{k=0}^{K} \pi_k^{old} p(\mathbf{y}_m | n_m; \gamma_k^{old})}.$$

M: Find γ^{new} and π^{new}, the values of γ and π, respectively, that maximize

$$\begin{aligned} Q(\gamma, \pi) &\equiv \sum_{m=1}^{M} \sum_{k=0}^{K} \hat{z}_{km} \log(\pi_k p(\mathbf{y}_m | n_m; \gamma_k)) \\ &= \sum_{m=1}^{M} \sum_{k=0}^{K} \hat{z}_{km} \log(\pi_k) + \sum_{m=1}^{M} \sum_{k=0}^{K} \hat{z}_{km} \log p(\mathbf{y}_m | n_m; \gamma_k) \quad (6) \end{aligned}$$

subject to constraint $\sum_{k=0}^{K} \pi_k = 1$.

Maximizing the first quantity in (6) via Lagrangian optimization gives

$$\pi^{new} = \frac{1}{M} \sum_{m=1}^{M} \hat{\mathbf{z}}_m,$$

where $\hat{\mathbf{z}}_m = (\hat{z}_{0m}, \hat{z}_{1m}, ..., \hat{z}_{Km})^T$. To get γ^{new}, note that the second quantity in (6) is proportional to

$$\sum_{k=1}^{K} g_k(\gamma_k) \equiv \sum_{k=1}^{K} \sum_{m=1}^{M} \hat{z}_{km} \sum_{n=1}^{N} y_{nm} \left[\gamma_k x_n - \log \left(\sum_{n=1}^{N} \exp\{\gamma_k x_n\} \right) \right] \quad (7)$$

The expression in (7) can now be maximized using any standard optimization method, such as a Newton-Rhapson routine. In fact it is important to note that a one-dimensional optimization routine can be applied K times to maximize each $g_k(\gamma_k)$, i.e. multi-dimensional optimization can be avoided. Now, the E and M steps can be repeated until, say, $l(\gamma^{new}, \pi^{new}) - l(\gamma^{old}, \pi^{old}) < \epsilon$, yielding maximum likelihood estimators denoted $\hat{\gamma}$ and $\hat{\pi}$.

6 The thresholding effect

This section demonstrates that conditional lFDR procedures tend to discover more non-negligible effects and fewer negligible effects than the usual (unconditional) lFDR procedures. The main result is first motivated using the data in Table 1.

To facilitate a simple and broadly applicable comparison of clFDR and lFDR procedures, we consider model (5) with two components, i.e. $\beta_m = 0$ or $\beta_m = \gamma_1 > 0$. We also utilize a normal approximation for each test statistic $T_m = \mathbf{x}^T \mathbf{Y}_m$. Recall that for $\beta_m = \gamma_1 > 0$ and conditionally upon $N_m = n_m$, the Z-score for T_m has mean $\mu(n_m, \gamma_1)$ and variance $\sigma^2(\gamma_1)$. See expressions (3) and (4), respectively. Thus, by the central limit theorem and delta method, given $N_m = n_m$ and $\beta_m = \gamma_1$, Z_m is asymptotically normal (as $n_m \to \infty$) with mean $\mu(n_m, \gamma_1)$ and variance $\sigma^2(\gamma_1)$. Here, $\theta_m = I(\beta_m = \gamma_1)$ and again denote mixing proportions by π_0 and π_1. In this section we suppress γ and π in the notation when possible for brevity.

The clFDR and lFDR procedures for comparison are based on normal approximations for the conditional and marginal mixture densities for Z-scores above. Specifically, denote the conditional mixture density by

$$f(z|n) = \pi_0 \phi(z; 0, 1) + \pi_1 \phi(z; \mu(n, \gamma_1), \sigma^2(\gamma_1)), \tag{8}$$

where $\phi(\cdot; a, b)$ denotes a normal density function with mean a and variance b. Define

$$clFDR(z, n) = \frac{\pi_0 \phi(z; 0, 1)}{f(z|n)}.$$

Assume that $N_1, N_2, ..., N_M \overset{i.i.d.}{\sim} p(n)$ and denote the support of N_m by $\mathcal{N}$. Then the marginal mixture density for (8) is

$$f(z) = \pi_0 \phi(z; 0, 1) + \pi_1 \sum_{n \in \mathcal{N}} \phi(z; \mu(n, \gamma_1), \sigma^2(\gamma_1)) p(n) \tag{9}$$

and the (marginal) local FDR is

$$lFDR(z) = \frac{\pi_0 \phi(z; 0, 1)}{f(z)}.$$

Observe that $f(z)$ and $lFDR(z)$ necessarily depend upon $p(n)$. In this section we use the empirical probability mass function depicted in Figure 1 for $p(n)$ in all computations, unless otherwise specified.

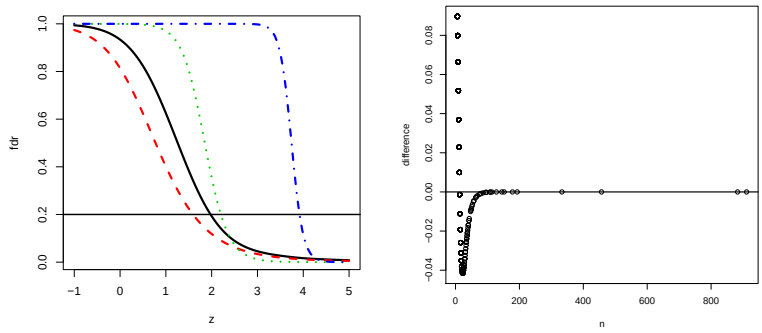

Figure 2: Left: $clFDR(z, n)$ vs. z for $n = 5$ (- - -), $n = 25$ ($\cdot\ \cdot\ \cdot$), and $n = 100$ (- $\cdot$ -) and $lFDR(z)$(—) vs. z for $\gamma_1 = 1$ and $\pi_0 = 0.5$. Right: difference in correct rejection probabilities (clFDR - lFDR) vs. n_m for n_m in Table 1.

The left panel in Figure 2 plots $lFDR(z)$ and $clFDR(z, n)$ vs. z for $\pi_0 = 0.5$, $\gamma_1 = 1$, $n = 5, 25, 100$. Recall that lFDR and clFDR procedures reject H_m if $clFDR(z_m, n_m) \leq \lambda$ or if $lFDR(z_m) \leq \lambda$ for some cutoff λ. For sake of illustration, we take $\lambda = 0.2$ in what follows. Observe in Figure 2 that $lFDR(z) \leq 0.2$ if $z > 1.98$ regardless of n. See the Appendix for exact expressions and discussion regarding the approximate equivalence of these rejection regions. The important point is that even though the rejection regions do not depend on n_m, the probability of correctly rejecting H_m (the power) is large when n_m is large and small when n_m is small due to the fact that $\mu(n_m, \gamma_1)$ increasing in n_m. For example, if $n_m = 5$ then under model (8) Z_m has a normal distribution with mean $\mu(5, 1) = 1.59$ and variance $\sigma^2(1) = 0.89^2$ when H_m is false. Thus the power is $1 - \Phi(1.98; 1.59, 0.89^2) = 0.33$, where $\Phi(\cdot; a, b)$ is the normal cumulative distribution function with mean a and variance b. When $n_m = 25$, $\mu(25, 1)$ is 3.55 and the power is 0.96. When $n_m = 100$, the power is 1.00. This same phenomenon occurs for any $\gamma_1 > 0$, including values of γ_1 that are only negligibly different from 0.

Procedures based on the conditional lFDR dampen this effect by decreasing the threshold for rejection when n_m is small and increasing it when n_m is large. For example, when $n_m = 5$, $clFDR(z_m, n_m) \leq 0.2$ when $z_m \geq 1.6$. That is, the Z-score threshold for rejection based on the clFDR is smaller than the Z-score threshold based on the lFDR. Consequently, the power of the test based on the clFDR is larger in this small n_m setting; it is 0.5 rather than 0.41. Now suppose that $n_m = 25$. Then $clFDR(z_m, n_m) \leq 0.2$ if $z_m \geq 2.2$. Because the Z-score threshold for the clFDR procedure is now larger than the threshold for the lFDR procedure, the power for the clFDR procedure is smaller: it is 0.94 rather than 0.98. In summary, conditioning on the N_ms in the local FDR analysis ensures that Z-score thresholds for rejection increase in n_m. Consequently, the

probability of discovering a negligible effect due to n_m being large is decreased and the probability of discovering a non-negligible effect when n_m is small is increased. See the right panel of Figure 2.

Theorem 2 states that the above phenomenon, where thresholds for rejection increase as n_m increases, occurs for large enough n_m as desired. See the Appendix for the proof and a discussion regarding the implied approximation of the rejection region $[clFDR(z,n) \leq \lambda]$ with $[z \geq a(n)]$.

Theorem 2 *Let $a(n)$ denote the smallest solution to $clFDR(z,n) = \lambda$. Under model (8), $a(n)$ is increasing in n whenever*

$$\mu(n,\gamma_1)^2 > 2\log\left(\sigma(\gamma_1)\frac{\pi_0(1-\lambda)}{(1-\pi_0)\lambda}\right). \tag{10}$$

Lets consider some specific settings. First recall that $\lambda = clFDR_{(k)}$ satisfies $\frac{1}{k}\sum_{i=1}^{k} clFDR_{(i)} \leq \alpha$. That is the average clFDR among the k rejected null hypothesis is near α. Hence, the threshold $\lambda = clFDR_{(k)}$ should be larger than α, especially if the clFDR statistics have a skewed right distribution. Hence, we consider $\lambda = 0.1$ and $\lambda = 0.2$. We also consider $\pi_0 = 0.5$ and 0.8 in our discussion. Each line in Figure 3 represents values of n and γ_1 where the inequality in (10) is an equality. The inequality is satisfied whenever (n, γ_1) is to the right of a line. For example, when $\lambda = 0.2$ and $\pi_0 = 0.5$, we see that the inequality is satisfied for every $n \geq 10$ if $\gamma_1 = 1$. If $\gamma_1 = 0.5$ instead of 1, then the inequality is satisfied for n greater than or equal to 25. When $\lambda = 0.1, \pi_0 = 0.5$, and $\gamma_1 = 0.5$ so that tests are generally less powerful, the inequality is satisfied for $n \geq 35$. The fact that the thresholds don't begin to increase until n is moderate in low power settings (smaller γ and λ) is to be anticipated as the power of the test is still relatively small when n is moderate. Hence, the probability of discovering a negligible effect is still small for moderate n.

7 Application

Here we analyze the data in Table 1 using the adaptive clFDR procedure in Section 5 and compare the results to the adaptive lFDR procedure results for the Z-scores in Section 3. Both procedures are applied at $\alpha = 0.05$. For the adaptive clFDR procedure, we consider a 3 and 4 component mixture model and use convergence criterion $l(\gamma^{new}, \pi^{new}) - l(\gamma^{old}, \pi^{old}) < 10^{-8}$. The EM algorithm failed to converge after 1000 iterations for a 5 component mixture model. Parameter summaries and the number of discoveries are in Table 3. R code for implementing the EM algorithm is available upon request.

Observe that all non-null component probability mass functions signify practical significance regardless of the number of components considered. That is, $\hat{\gamma}_k$ is more than arbitrarily different from 0. For example, $\mathbf{p}(0.78) = [0.08, 0.11, 0.16, 0.25, 0.40]$ certainly deviates from $\mathbf{p}(0) = \mathbf{0.2}$ by an amount that would signify a non-negligible association with productivity.

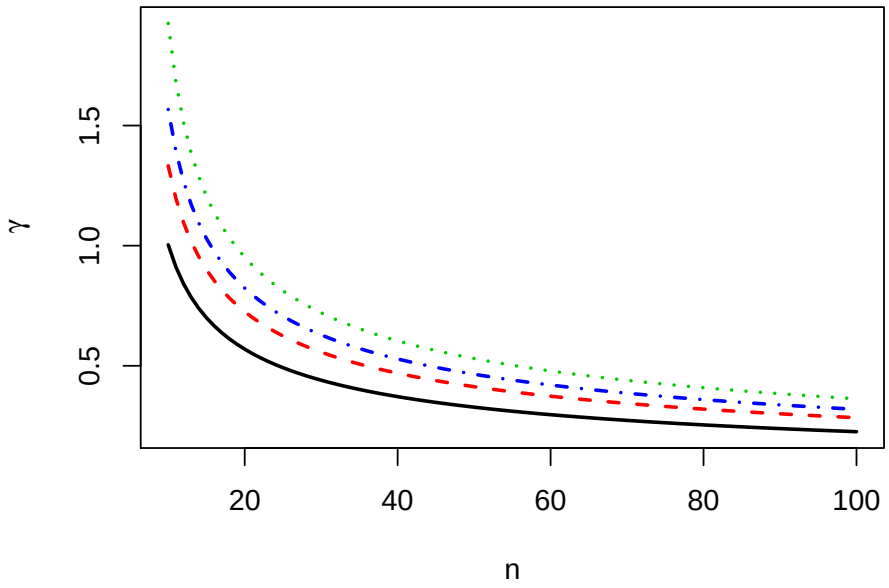

Figure 3: Depiction of (n, γ_1) values satisfying equation (10). Lines represent values of n and γ_1 where the inequality in (10) is an equality for $(\lambda, \pi_0) = (0.1, 0.8), (0.2, 0.8), (0.1, 0.5), (0.2, 0.5)$ from top to bottom. The inequality is satisfied whenever (n, γ_1) is to the right of the line.

Table 3: Parameter estimates, AIC, BIC and the number of discoveries for Models 1 and 2.

Model	$\hat{\pi}_0$	$(\hat{\pi}_1, \hat{\gamma}_1)$	$(\hat{\pi}_2, \hat{\gamma}_2)$	$(\hat{\pi}_3, \hat{\gamma}_3)$	AIC	BIC	R($\mathbf{Y}$)
1	0.69	(0.16, -1.13)	(0.15,0.78)	NA	12221.9	12240.5	99
2	0.69	(0.03, -2.68)	(0.13, -1.03)	(0.15, 0.79)	12145.6	12173.5	97

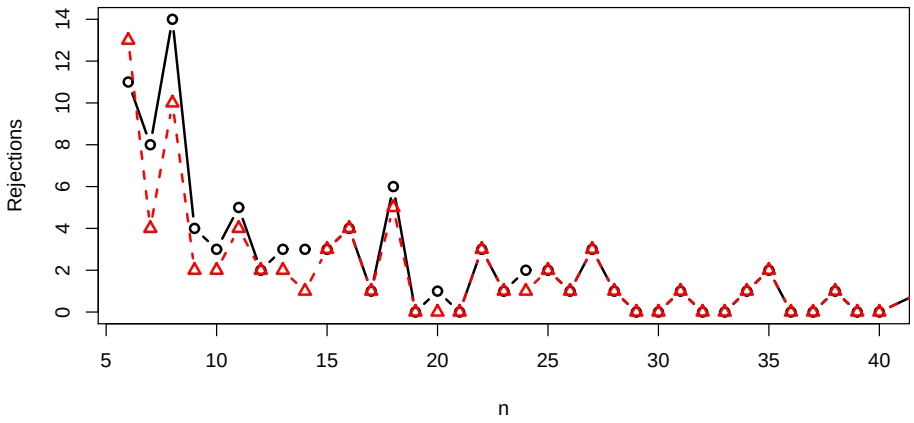

Figure 4: The number of rejections vs. n for the conditional lFDR procedure (o) and for the lFDR procedure (△).

A more careful inspection of the results confirms that the clFDR procedure tends to discover more non-negligible effects and fewer negligible effects. Figure 4 plots the number of discoveries vs. n for the 3 component normal mixture model in Table 2 and for the 3 component multinomial mixture model in Table 3. Observe that the clFDR procedure makes more discoveries when n_m is small, as Theorem 2 suggests. Also, when $n_m \geq 40$ (not plotted), the clFDR procedure makes two fewer discoveries. These latter discrepancies occur for $n_m = 193$ and $n_m = 911$.

In general, the characteristics of the data when discrepancies exist are as anticipated. For example, $\mathbf{Y}_m^T/n_m = (5, 7, 0, 1, 1)/11 = (0.56, 0.64, 0, 0.09, 0.09)$ is certainly more than negligibly different from $\mathbf{p}(0) = \mathbf{0.2}$, but was declared significant only by the clFDR procedure. Here n_m is small. On the other hand, $\mathbf{Y}_m^T/n_m = (134, 117, 252, 231, 177)/911 = (0.15, 0.13, 0.28, 0.25, 0.19)$ deviates from $\mathbf{p}(0)$ less severely but n_m is large. Consequently, it was discovered by the lFDR procedure but not by the clFDR procedure.

Appendix

First we verify that for $\sigma^2(\gamma_1) \neq 1$, the rejection region $\{z : clFDR(z, n) \leq \lambda\}$ can be written $\{z : a(n) \leq z \leq b(n)\}$ and then demonstrate that it is well approximated by $\{z : z \geq a(n)\}$ under model (8). Expressions for $a(n)$ and $b(n)$ are also derived. For brevity, denote $\mu(n, \gamma_1)$ by $\mu(n)$ and $\sigma^2(\gamma_1)$ by σ^2.

Observe that setting $clFDR(z, n) = \lambda$, plugging expressions for the normal

densities in $clFDR(z,n)$, and rearranging we get

$$z^2(\sigma^2-1)+2z\mu(n)-2\sigma^2 log(\sigma k)-\mu(n)^2=0,$$

where $k=\frac{\pi_0(1-\lambda)}{(1-\pi_0)\lambda}$. The solutions of this quadratic ($a(n)$ and $b(n)$) are

$$\frac{\mu(n)\pm\sqrt{\mu(n)^2\sigma^2-2(1-\sigma^2)\sigma^2\log(\sigma k)}}{1-\sigma^2}.$$

To verify that $\{z: clFDR(z,n)\leq\lambda\}=\{z: a(n)\leq z\leq b(n)\}$, first observe that taking the derivative of $clFDR(z,n)$ with respect to z gives

$$clFDR'(z,n)=\left[\frac{\pi_0(1-\pi_0)\phi(z;0,1)\phi(z;\mu(n),\sigma^2)}{f(z|n)^2}\right]\left[\frac{z(1-\sigma^2)-\mu(n)}{\sigma^2}\right]$$

which is negative for $z<\mu(n)/(1-\sigma^2)$ and positive for $z>\mu(n)/(1-\sigma^2)$. Hence, $clFDR(z,n)$ is "U" shaped and hence $clFDR(z,n)\leq\lambda$ if and only if $a(n)\leq z\leq b(n)$.

It should be noted that if $\sigma^2=1$ then $clFDR'(z,n)\leq 0$ for $z\in\Re$ and the rejection region is of the form $[z>a]$ for some a. It can be verified, however, that $\sigma^2<1$ whenever $\gamma\neq 0$, thereby resulting in the above rejection region. This rejection rule is seemingly untractable because we fail to reject H_m if $Z_m>b(n)$ even though $Z_m>b(n)$ is stronger evidence against H_m than $a(n)\leq Z_m\leq b(n)$. However, the event that $Z>b(n)$ occurs with very small probability in practice and hence has little impact on the clFDR procedure. For example, $[b(n)-\mu(n)]/\sigma$ was at least 4.6 for all combinations of $\gamma=0.5,1,2;\lambda=0.05,0.1,0.2;n=5,100,1000$; and $\pi_0=0.1,0.5,0.9$ and is typically greater than 50. That is, $b(n)$ was at least 4.6 standard deviations above the mean when H_m was false, and hence the probability of such a Z-score is extremely small. The curious reader is referred to Cao et al. (2013) for a more detailed investigation.
Proof of Theorem 2: The derivative of $a(n)$ with respect to $\mu(n)$ is

$$\frac{1}{1-\sigma^2}\left[1-\frac{\mu(n)\sigma^2}{\sqrt{\mu(n)^2\sigma^2-2(1-\sigma^2)\sigma^2 log(\sigma k)}}\right]$$

which is positive when

$$\frac{\mu(n)\sigma^2}{\sqrt{\mu(n)^2\sigma^2-2(1-\sigma^2)\sigma^2 log(\sigma k)}}<1.$$

Some algebra gives that this inequality is satisfied if and only if $\mu(n)^2>2\log(\sigma k)$. Plugging in $\pi_0(1-\lambda)/[(1-\pi_0)\lambda]$ for k gives the desired inequality.

References

Alami, Y., W. Achouak, C. Marol, and T. Heulin (2000). Rhizosphere soil aggregation and plant growth promotion of sunflowers by an exopolysaccharide-producing rhizobium sp. strain isolated from sunflower roots. *Applied and Environmental Microbiology 66*(8), 3393 – 3398.

Anderson, M. and J. Habiger (2012). Characterization and identification of productivity-associated rhizobacteria in wheat. *Applied and Environmental Microbiology 78*(12), 4434 – 444.

Bakker, P., R. Doornbos, C. Zamioudis, R. L. Berendsen, and C. M. J. Pieterse (2013). Induced systemic resistance and the rhizosphere microbiome. *Plant Pathology Journal 29*(2), 136 – 143.

Benjamini, Y. and Y. Hochberg (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B. Methodological 57*(1), 289–300.

Blanchard, G. and E. Roquain (2009, December). Adaptive FDR control under independence and dependence. *Journal of Machine Learning Research 10*, 2837 –2831.

Cai, T. T. and W. Sun (2009). Simultaneous testing of grouped hypotheses: finding needles in multiple haystacks. *Journal of the American Statististical Association 104*(488), 1467–1481.

Cao, H., W. Sun, and M. R. Kosorok (2013). The optimal power puzzle: scrutiny of the monotone likelihood ratio assumption in multiple testing. *Biometrika 100*(2), 495–502.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B. Methodological 39*(1), 1–38. With discussion.

Efron, B. (2010). *Large-Scale Inference*. Cambridge: Cambridge University Press.

Efron, B., R. Tibshirani, J. D. Storey, and V. Tusher (2001). Empirical Bayes analysis of a microarray experiment. *Journal of the American Statistical Association 96*(456), 1151–1160.

Gans, J., M. Wolinsky, and J. Dunbar (2005). Computational improvements reveal great bacterial diversity and high metal toxicity in soil. *Science 309*, 1387–1390.

Genovese, C. and L. Wasserman (2002). Operating characteristic and extensions of the false discovery rate procedure. *Journal of the Royal Statistical Society, Series B 64*(3), 499–517.

Liang, K. and D. Nettleton (2012). Adaptive and dynamic adaptive procedures for false discovery rate control and estimation. *Journal of the Royal Statistical Society. Series B. 74*(1), 163–182.

McCullagh, P. and J. A. Nelder (1989). *Generalized linear models*. Monographs on Statistics and Applied Probability. Chapman & Hall, London.

Mendes, R., P. Garbeva, and J. Raaijmakers (2013). The rhizosphere microbiome: significance of plant beneficial, plant pathogenic, and human pathogenic microorganisms. *Microbiology Reviews 37*(5), 634–663.

Storey, J. (2003). The positive false discovery rate: a bayesian interpretation and the q-value. *Annals of Statistics 31*(6), 2012 – 2035.

Storey, J. D., J. E. Taylor, and D. Siegmund (2004). Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society. Series B 66*(1), 187–205.

Sun, W. and T. T. Cai (2007). Oracle and adaptive compound decision rules for false discovery rate control. *Journal of the American Statistical Association 102*(479), 901–912.

Sun, W. and A. C. McLain (2012). Multiple testing of composite null hypotheses in heteroscedastic models. *Journal of the American Statistical Association 107*(498), 673–687.

Suslow, T. and M. Schroth (1982). Role of deleterious rhizobacteria as minor pathogens in reducing crop growth. *Phytopathology 72*(1), 111 – 115.

Vessey, J. K. (2003). Plant growth promoting rhizobacteria as biofertilizers. *Plant and Soil 255*(2), 571 – 586.

Yang, J., J. Kloepper, and C. Ryu (2009). Rhizosphere bacteria help plants tolerate abiotic stress. *Trends in Plant Science 14*(1), 1 – 4.